



read-2suv-1-3-compact-30.loop : 1.8.1-937-rel-enterprise : 2.0.0-1495-rel-enterprise




(righter is better)


	1.8.1 – 937	2.0.0 – 1495
<i>Runtime (in hr)</i>	6	5.48
<i>Avg. Drain Rate</i>	294.06	216.15
<i>Peak Disk (GB)</i>	93.15	129.43
<i>Peak Memory (GB)</i>	17.06	16.37
<i>Avg. OPS</i>	12.66K	15.01K
<i>Avg. mem memcached (GB)</i>	16.98	15.97
<i>Avg. mem beam.smp (MB)</i>	71.18	251.6
<i>Latency-get (90th) (ms)</i>	1.49	0.86
<i>Latency-get (95th) (ms)</i>	2.59	1.76
<i>Latency-get (99th) (ms)</i>	4.74	3.67
<i>Latency-set (90th) (ms)</i>	1.73	1.13
<i>Latency-set (95th) (ms)</i>	2.74	2.05
<i>Latency-set (99th) (ms)</i>	4.8	3.68
<i>Latency-query (80th) (ms)</i>	NA	NA
<i>Latency-query (90th) (ms)</i>	NA	NA
<i>Latency-query (95th) (ms)</i>	NA	NA
<i>Latency-query (99th) (ms)</i>	NA	NA
<i>Latency-query (99.9th) (ms)</i>	NA	NA
<i>Avg. QPS</i>	0	0
<i>Rebalance Time (sec)</i>	0	0
<i>Testrunner Version</i>	ae1ad17	cab63e8



ep queue size



ep_diskqueue_drain




buildinfo.version


- 1.8.1-937-rel-enterprise
- 2.0.0-1495-rel-enterprise

Total drained items on disk queue


ep_bg_fetched ops/sec



tmp_oom ops/sec




vb_active_eject/sec




Number of times item values got ejected

vb_replica_eject/sec



ep_tap_replica_queue_backoff/sec




buildinfo.version


- 1.8.1-937-rel-enterprise
- 2.0.0-1495-rel-enterprise

*Number of back-offs received per second
while sending data over replication
TAP connections*


vb_active_resident_items_ratio




vb_replica_resident_items_ratio



curr_items




cur_items_total




Total number of items


mem_used



Docs data size




Docs disk size




- buildinfo.version**
- 1.8.1-937-rel-enterprise
 - 2.0.0-1495-rel-enterprise

The size of all data files for this bucket, including the data itself, meta data and temporary files (2.0 only)


Docs actual disk size




Views data size




Views disk size




Views actual disk size




Total disk size



Docs fragmentation




Views fragmentation




cmd_get ops/sec




cmd_set ops/sec




of get misses



of get hits




cache_miss percentage




cache_miss percentage 0-5



Number of connections




CPU utilization – 10.2.1.65:8091




buildinfo.version
● 1.8.1-937-rel-enterprise
● 2.0.0-1495-rel-enterprise

CPU utilization – 10.2.1.66:8091




buildinfo.version
● 1.8.1-937-rel-enterprise
● 2.0.0-1495-rel-enterprise


CPU utilization – 10.2.1.67:8091



CPU utilization – 10.2.1.68:8091



cpu_time_diff: memcached – 10.2.1.65



buildinfo.version


- 1.8.1-937-rel-enterprise
- 2.0.0-1495-rel-enterprise

*CPU tick differences (jiffies)
among snapshots across time*


cpu_time_diff : beam.smp - 10.2.1.65



cpu_time_diff: memcached – 10.2.1.66



cpu_time_diff : beam.smp - 10.2.1.66




buildinfo.version

- 1.8.1-937-rel-enterprise
- 2.0.0-1495-rel-enterprise


*CPU tick differences (jiffies)
among snapshots across time*

----time (sec)---->


cpu_time_diff: memcached – 10.2.1.67




cpu_time_diff : beam.smp - 10.2.1.67




cpu_time_diff: memcached – 10.2.1.68




cpu_time_diff : beam.smp - 10.2.1.68



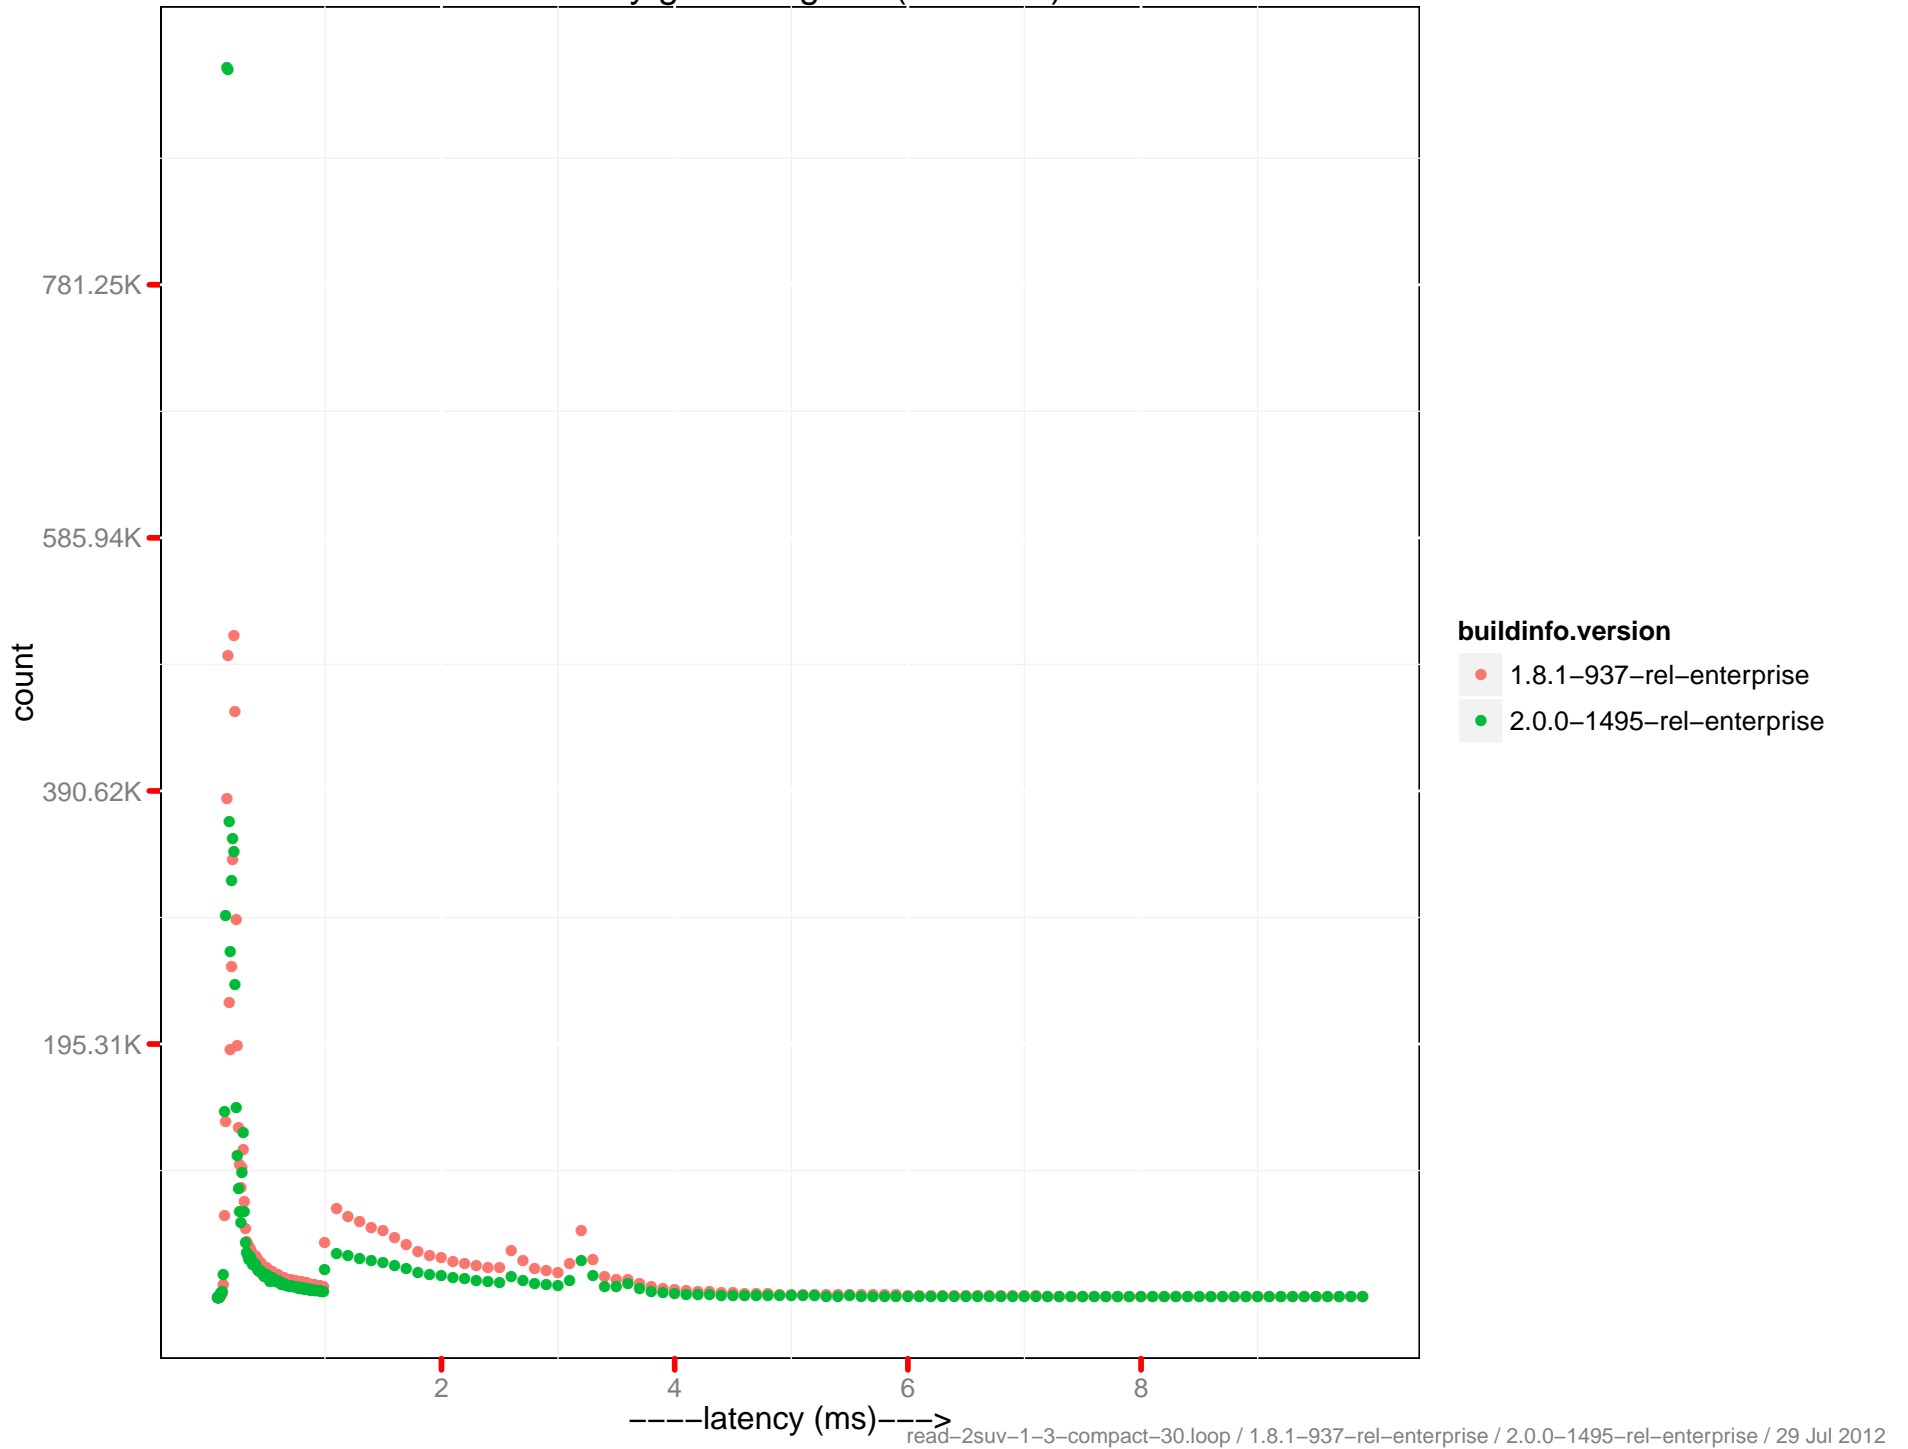
Data disk size




Latency get histogram




Latency get histogram (0-10 ms)




Latency set histogram




Latency set histogram (0–10 ms)




Latency-get 90th percentile




Latency-get 90th percentile (0 - 10ms)



Latency-get 95th percentile




Latency-get 95th percentile (0 - 10ms)




buildinfo.version
● 1.8.1-937-rel-enterprise
● 2.0.0-1495-rel-enterprise


Latency-get 99th percentile




Latency-get 99th percentile (0 - 10ms)




Latency-set 90th percentile




Latency-set 90th percentile (0 - 10ms)




Latency-set 95th percentile




Latency-set 95th percentile (0 - 10ms)




Latency-set 99th percentile



Latency-set 99th percentile (0 - 10ms)



Query throughput



```
read-2suv-1-3-compact-30.conf
# read 20M load, 2M hot reload, 9M access creates, draining
# speed limit = 3k per node
#
# system memory: 20G per node

performance.eperf.EPerfClient.test_eperf_read

params:

# general
batch=50
kind=nonjson
mem_quota=20000
spec=read-2suv-1-3-compact-30
db_compaction=30

# load phase
hot_init_items=2000000
items=20000000

# access phase
# Read:Insert:Update:Delete Ratio = 90:3:6:1.
ratio_sets=0.1
ratio_misses=0.05
ratio_creates=0.30
ratio_deletes=0.1428
ratio_hot=0.05
ratio_hot_gets=0.99
ratio_hot_sets=0.99
ratio_expirations=0.005
max_creates=9000000

# control (defaults: pytest/performance/perf_defaults.py)
load_wait_until_drained=1
loop_wait_until_drained=0
mcsoda_heartbeat=3
mcsoda_max_ops_sec=500
tear_down=1
tear_down_proxy=1
tear_down_bucket=0
tear_down_cluster=1
tear_down_on_setup=0
```

```
vesta.ini
[global]
username:root
password:couchbase
port:8091
data_path:/data

[servers]
1:10.2.1.65
2:10.2.1.66
3:10.2.1.67
4:10.2.1.68

[clients]
1:10.2.1.60

[membase]
rest_username:Administrator
rest_password:password

[dashboard]
1:dashboard.hq.couchbase.com:80
```